

# ***Challenges with Data Integrity***

***Henry Newman***

***Instrumental, Inc.***

***hsn@instrumental.com***

# *What is covered*

- Problem Statement
- Flash concerns
- The numbers

# ***Problem Statement***

***Significant changes are needed  
in error encoding***

- Fibre channel originally developed at 25 MB/sec and it is now 64 times faster at 1600 MB/sec
- IDE channel originally was .625 MB/sec and it is now 480 times faster at 300 MB/sec
- The channel error rate for both is  $10E^{12}$  bits
  - Corrected to higher value
  - IB copper is the same
  - Optical is claimed to be better
  - 10GbE is the same

- Research has shown rates of errors undetected by link CRC's and TCP checksums ranging from one in 16 million to 10 billion packets
  - <http://portal.acm.org/citation.cfm?doid=347059.347561>
  - 16 Million 9K packets can be sent in less than 2 minutes on a 10GE link
  - 10 Billion 9K packets can be sent in less than 24 hours on a 10GE link
  - “When compared to un-detected error rates for local I/O (e.g., disk drives), these rates are disturbing”
- This is with TCP/IP checksums!

- Both disks and tape have far more error encoding than the channels
  - The encoding for tape is far more robust; for example, LTO is around 8 orders of magnitude greater than FC
  - Enterprise tape is at least 10 orders of magnitude better
- The channel error encode was not a consideration 20 years ago as things were too slow and too expensive to have lots of channels
  - No longer true

- Robustness of error encoding has not changed for either storage channel type over the life of the channel
  - Between 20 (FC) and 25+ (SATA) years
  - It needs to be changed in ethernet
  - Seagate recently published SAS/FC undetectable rates
- This has resulted in a situation where organizations are starting to see actual data loss as we have hit the wall with error encoding
  - This does not even consider PCIe, memory issues or other part of the path

- Everyone thinks these are the ultimate solution for metadata and logs
- SSD have wide performance range for read and especially write
- NAND flash does not support writes over 100K times to a specific location and flash will fail
  - What happens to reliability at 70K, 90K and just before failure of the write?
  - What historical data do we have?

- SMART is a standard that was developed for disk drives
  - Some of the error conditions found in flash do not fit within the framework for SMART
  - It took RAID vendors 3-5 years to accomplish predictive failure in controllers for disk drives
- No standard for SMART statistics for flash
  - New proposal to ANSI for flash but in early stages

		Annual Failure Rates at Different Sustained Transfer Rates Per Second.						
Est.T10 PI Detection	UDBER	0.5 GB/sec	1 GB/sec	10 GB/sec	100 GB/sec	1 TB/sec	10 TB/sec	100 TB/sec
	1.E-28	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.E-27	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.E-26	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.E-25	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.E-24	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.E-23	0.0	0.0	0.0	0.0	0.0	0.0	0.3
	1.E-22	0.0	0.0	0.0	0.0	0.0	0.3	2.7
FC/SAS	1.E-21	0.0	0.0	0.0	0.0	0.3	2.7	27.1
	1.E-20	0.0	0.0	0.0	0.3	2.7	27.1	270.9
	1.E-19	0.0	0.0	0.3	2.7	27.1	270.9	2708.9
	1.E-18	0.1	0.3	2.7	27.1	270.9	2708.9	27089.2
SATA	1.E-17	1.4	2.7	27.1	270.9	2708.9	27089.2	270892.2
	1.E-16	13.5	27.1	270.9	2708.9	27089.2	270892.2	2708921.8
	1.E-15	135.4	270.9	2708.9	27089.2	270892.2	2708921.8	27089217.7

- These annual failure rates are for a perfect world where
- the channels are operating at the specified rate of  $10E^{-12}$  and corrected to  $10E^{-17/19}$
- What happens when the world is not perfect?
- What about GbE?

Technology	Unrecoverable read error per bits read	1 PB	10 PB	40 PB	100 PB
1 TB Consumer SATA	10E14	9.007	90.07	360.288	900.720
1 TB	10E15	0.901	9.007	36.029	90.072
450 GB	10E16	0.090	0.901	3.603	9.007
LTO-4/TS1130	10E17	0.009	0.090	0.360	0.901
T10000B	10E19	0.000	0.001	0.004	0.009

- Clearly this is a problem that needs to be addressed
  - Vendors do not seem to be improving these values as it is on required in the commodity world
- To ensure data reliability other methods need to be investigated

		1 PB	10 PB	40 PB	100 PB
Network	Data Rate Gb/sec	Days to Replicate			
OC-3	0.15	802	8018	32071	80178
OC-12	0.61	200	1998	7992	19980
OC-48	2.40	51	506	2023	5057
OC-192	9.60	13	126	506	1264
OC-384	19.20	6	63	253	632
OC-768	38.40	3	32	126	316

- Given hard error rates and time to replicate in the event of a disaster at a site data will be lost
- I do not believe that “Hadoop method” will work given these considerations especially with the cost of power for CPUs and memory
  - 5 year costs with power is huge compared with other methods and risk of data loss in case of disaster is an issue

- If there is corruption most people blame the file system first and the hardware last
  - That might have been a good plan in the 1970s-1990s but it is no longer true in most cases
- Some questions we could discuss as I have some thoughts and opinions:
  - Does error correction belong in the file system?
  - What should be done about hard error rate?
  - What will happen to tape given Dedup impact?